



HAL
open science

Exploring the Placement and Design of Word-Scale Visualizations

Pascal Goffin, Wesley Willett, Jean-Daniel Fekete, Petra Isenberg

► **To cite this version:**

Pascal Goffin, Wesley Willett, Jean-Daniel Fekete, Petra Isenberg. Exploring the Placement and Design of Word-Scale Visualizations. IEEE Transactions on Visualization and Computer Graphics, Institute of Electrical and Electronics Engineers, 2014, 20 (12), pp.2291–2300. 10.1109/TVCG.2014.2346435 . hal-01024278

HAL Id: hal-01024278

<https://hal.inria.fr/hal-01024278>

Submitted on 17 Dec 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Exploring the Placement and Design of Word-Scale Visualizations

Pascal Goffin, Wesley Willett, Jean-Daniel Fekete *Senior Member, IEEE* and Petra Isenberg

Abstract—We present an exploration and a design space that characterize the usage and placement of word-scale visualizations within text documents. Word-scale visualizations are a more general version of sparklines—small, word-sized data graphics that allow meta-information to be visually presented in-line with document text. In accordance with Edward Tufte’s definition, sparklines are traditionally placed directly before or after words in the text. We describe alternative placements that permit a wider range of word-scale graphics and more flexible integration with text layouts. These alternative placements include positioning visualizations between lines, within additional vertical and horizontal space in the document, and as interactive overlays on top of the text. Each strategy changes the dimensions of the space available to display the visualizations, as well as the degree to which the text must be adjusted or reflowed to accommodate them. We provide an illustrated design space of placement options for word-scale visualizations and identify six important variables that control the placement of the graphics and the level of disruption of the source text. We also contribute a quantitative analysis that highlights the effect of different placements on readability and text disruption. Finally, we use this analysis to propose guidelines to support the design and placement of word-scale visualizations.

Index Terms—Information visualization, text visualization, sparklines, glyphs, design space, word-scale visualizations

1 INTRODUCTION

Small high-resolution data graphics, included alongside words or word sequences in text documents, can often communicate information that could not be succinctly conveyed by the text itself. Examples include small stock charts embedded next to the name of a company, game statistics next to the name of a soccer team, or weather trends next to the name of a city. Traditionally, most of these “word-scale visualizations” have consisted of small line charts and bar charts and been placed in-line with text. Edward Tufte terms these word-scale visualizations “sparklines” [30], and provides some guidelines for their visual design. However, Tufte provides little guidance for *placing* word-scale visualizations with respect to text, suggesting only that they be placed in a “relevant context”—usually just after the word that they complement. However, the space of design and placement options for word-scale visualizations is actually quite large, and the consequences of placement decisions, in particular, are not well-understood.

In this paper, we provide design considerations for placing word-scale visualizations associated with words or word sequences (what we refer to as “entities”) in a document. Our work is motivated by a close collaboration on digital note-taking with historians in the digital humanities. When visiting an archive, the historians we work with regularly take detailed notes on their findings. In these notes, they specifically tag entities such as the people, locations, or dates that occur in their document sources. The goal of tagging these entities is to help historians build an understanding of how entities relate to one another, where else the same entities appear in their notes, and what kinds of metadata are associated with them. Embedding this information using word-scale visualizations is a promising approach, because these small visualizations can add additional information in-context without distracting attention from the primary reading task.

In prior work, sparklines have typically been placed before or after the word they are related to. However, this is often not possible for the kinds of notes taken by our historians—e. g. when adding information to scanned documents and other immutable texts. Placing word-scale visualizations in-line with text may also be undesirable in other situations, as it requires reflowing the text and restricts the visu-

alization’s maximum height to that of the font—making visualizations hard to read when small font sizes were chosen. In-line visualizations can also disrupt sentences, making the text more difficult to read.

To better understand the options available for integrating word-scale visualizations in text documents, we outline a design space of possible placements relative to the text. In doing so, we relax some aspects of Tufte’s original sparkline definition, imposing less restrictive size requirements and allowing the small visualizations to extend beyond strictly “word sized.” Also, while Tufte did not restrict sparklines to specific visual encodings, the term “sparkline” does inherently suggest a “line-based” data encoding such as a line chart. In contrast, we specifically allow a variety of encodings, including geographical maps, heat maps, pie charts, and more complex visualizations and, thus, chose the term *word-scale visualizations*. We also formalize the notion of an *entity*—a concrete piece of text with associated meta-data that can be encoded in a word-scale visualization. This explicit connection between an entity and a word-scale visualization directly affects the options for placing the visualization, and allows us to formally characterize the spatial relationship between text and graphic.

We begin our discussion by reviewing related work on small-scale and text visualizations. Then, in Section 3 we introduce the design space, its focus, and dimensions. Section 4 details several placement options and discusses trade-offs between word-scale visualization placement options. In Section 5 we discuss three examples that demonstrate the importance of the association between word-scale visualization and entity for the purpose of layout and interaction. Finally, in Section 6 we provide an in-depth analysis that examines how various placement options affect word-scale visualization placement in real documents. Based on this analysis, we provide recommendations that can help designers choose the right word-scale visualization given their own constraints.

2 RELATED WORK

Our work relates closely to four research areas: (a) the use of sparklines and the design of word-scale visualizations (b) the integration of meta-data within text documents, (c) research on labeling in visualization, and (d) the readability of texts and visualizations.

2.1 Sparklines and Small-Scale Visualizations

According to Tufte [30] sparklines are “small, intense, simple, word-sized graphics with typographic resolution” that can be included anywhere a word or number can be—e. g. in a sentence, table, headline, map, spreadsheet or graphic. Tufte presents several examples of these embeddings. One example shows sparklines embedded in-line with text in order to provide metadata for a single word, for example glucose measurements next to the word glucose. In another, sparklines

- Pascal Goffin is with Inria. E-mail: pascal.goffin@inria.fr.
- Wesley Willett is with Inria. E-mail: wesley.willett@inria.fr.
- Jean-Daniel Fekete is with Inria. E-mail: jean-daniel.fekete@inria.fr.
- Petra Isenberg is with Inria. E-mail: petra.isenberg@inria.fr.

Manuscript received 31 Mar. 2014; accepted 1 Aug. 2014; date of publication xx xxx 2014; date of current version xx xxx 2014.

For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org.

show temporal trends in a data table to provide information that would otherwise have required many columns.

Besides these examples, sparklines have been discussed relatively infrequently in the research literature. However, four recent approaches that meet our definition of word-scale visualizations include: SparkClouds [18], Gestaltlines [6], Separation Plots [14] and SportLines [20]. SparkClouds embed sparklines above words in a tag cloud to convey temporal trends. Gestaltlines were introduced by Brandes et al. as “gestalt theory-informed glyphs in sparklines.” The patterns that Gestaltlines show are designed to be perceived holistically and pre-attentively. Separation Plots show measurements that assess the predictive power of models with binary outcomes. They visually encode rows of data as cells in a small linear strip and overlay a line graph of the fitted values. SportLines are small word-sized visualizations that show sequences of passes in soccer games. They are different to traditional sparklines as they show a perspective map of a soccer field with a bold line showing the grouped passes. For Gestaltlines, Separation Plots and SportLines, the authors show embeddings in-line with text, similar to Tufte’s glucose example. In this paper, we show that other placement options exist, and can also be used to place Gestaltlines, Separation Plots, SportLines, and many other small-scale visualizations within text documents.

Glyphs are another type of small visualization related to sparklines but are not typically placed in the context of text documents. Instead, glyphs are often used as small-multiples or replace data marks in other visualizations, e. g. on maps, scatterplots, or in graphs. Yet, glyphs are also sometimes placed inside of text documents. Abdul-Rahman et al. [1], for example, use glyphs to encode and visualize a large set of poetic variables and place these glyphs in the space between lines.

An extensive overview of different general glyph types can be found in Borgo et al.’s [4] work, as well as in an overview paper by Ward [31]. In the context of this latter overview article [31], Ward also contributes a taxonomy of glyph placement strategies. He categorizes these strategies into data-driven and structure-driven approaches. A data-driven placement positions glyphs based on some of the dimensions of the represented data point, whereas structure-driven approaches use existing relationships between the data points (e. g. an order, hierarchy, or network links). The problem of placing a word-scale visualization next to an entity within a text document is related to the structure-driven placement option. Thus, Ward shares several placement concerns with our work, including overlap between glyphs and the addition of space to make room for glyphs.

Finally, there are a number of small visualizations besides sparklines and glyphs which share similar size constraints, but have not typically been embedded with text. Horizon graphs [11, 22, 24] are one such example. Horizon graphs use a split-space technique to represent a large number of time series in a single compact view with constrained vertical height. The height of the visualization can be kept very small by using an approach called “two-tone pseudo-coloring” [24]. Given their visual encoding, Horizon Graphs could certainly be used as a word-scale visualization to give meta-information about entities. Scented Widgets [32]—small data graphics embedded in GUI controls to support social navigation—showcase another opportunity for embedding small visualizations into other interfaces. Similarly, word-scale visualizations can be used to provide information scent cues within text documents—for example by showing visitation data for hyperlinks on web pages. We show an example of this kind of usage in Section 5.2.

In summary, while research has considered the problem of designing small visualizations and even embedding them in various contexts we are aware of no previous work that has explored the placement of word-scale visualizations in text documents.

2.2 Meta-Data within Text Documents

While the literature contains little advice on how to place small-scale visualizations in text, others have considered how to place general meta-data within text documents.

Annotations are a general category of meta-data placed in text documents. They can be as small as a single character but are sometimes

much more elaborate. In Japanese, for example, furiganas—small reading aids—are often placed on top of a word of interest to help translate kanjis (Chinese characters) into the Japanese alphabet.

Alternatively, some research has shown that the most preferred placement for larger annotations in a text document may be the margin [19]. Despite this finding, researchers have explored how to use and improve other placements. TextTearing [33], for example, allows readers to place annotations in the inter-line space via a lightweight interaction technique. Fluid Documents [7] similarly adds additional information into a text document by adjusting inter-line spacing. Fluid text includes links that expand in place to allow the text to still be read linearly while preserving the surrounding information. In an observational study Zellweger et al. [34] investigated different placement options for meta-information related to hypertext links (e. g. margins, footnotes, and in-line). They found a large variability in participants’ preferences but observed that, in general, it was desirable to keep meta-information close to the original link. Where information was placed also had a significant impact on reading speed with information close to the links being read significantly faster. An additional feature allowed meta-information to be “frozen” or kept visible in the view and participants used this feature regularly to compare different meta-information. This last feature is important for our design space and is one of the reasons we discuss both dynamic and static placements.

2.3 Labeling

The labeling problem is also related to our problem of placing word-scale visualizations in text documents. Bertini et al. [2] define labeling as attaching text labels to graphical marks to convey associated semantic information. Transferred to our problem, the entity is the graphical mark and the word-scale visualization is the label. Fekete and Plaisant [10] propose a taxonomy of labeling techniques with two categories: a) static labeling with the goal of finding a universal placement configuration for a whole visualization, and b) dynamic labeling where each label is treated as a dynamic object that can appear or disappear or change its location depending on the user’s interaction.

The transient visualization concept by Jakobsen and Hornbæk [17] relates to our work if we look at the labeling problem from the viewpoint of interaction and information-in-context. Their goal is to visualize information near the user’s focus of attention only when needed and without using extra screen space. This concept relates to our word-scale visualization placement problem since there are cases where a reader only wants to see the visualization when needed. We, thus, discuss both dynamic and static placement options.

2.4 Readability of Text and Small-Scale Visualization

Thus far, very little work has explored how word-scale visualizations might impact the readability of surrounding text. The most closely-related work is the previously mentioned study by Zellweger et al. [34] on Fluid Documents. The authors conducted eye tracking measurements and concluded that subjects were quite sensitive to different placement strategies but also that adjustments to document typography did not cause their point of focus to shift wildly.

Similarly, there is little existing research characterizing how small size affects the readability and perception of visualizations. For example, while several studies have compared the readability of different glyph designs (e. g. [12, 23]), we are aware of only one study (by Heer et al. [16]) that has explored the impact of size on readability. Heer et al. compared small line charts against two types of Horizon Graphs (1-band and 2-band) and investigated the impact of chart height on both designs. They found that small chart heights negatively impacted how accurately and quickly participants estimated the difference between two data points. The authors also saw that Horizon Graph design had an influence on the amount of estimation error and that both horizon graph designs performed better than line charts. The authors conclude with the recommendation to draw line charts of 24 pixels (6.8mm) height and 2-band Horizon Graphs at 6 or 12 pixels (1.7 / 2.4 mm)—all vertical heights that fit the definition of word-sized graphics at common font sizes and display resolutions.

When designing small-scale visualizations, the choice of aspect ratio impacts the perception of trends and patterns in the data. Cleveland et al. [8] recommend banking to 45° to choose the aspect ratio of a line chart. This is due to the fact that an average orientation of 45° maximizes the discriminability of adjacent line segments. Banking to 45° selects an aspect ratio for a chart such that the average orientation of all line segments is 45° . Based on this work, Heer and Agrawala [15] presented a multi-scale banking technique which identified the different trends in the data at various frequency scales to generate banked charts. Yet, Talbot et al. [28] showed that aspect ratio selection is not as simple as it may seem and that trade-offs between reducing errors for some tasks and increasing it for others have to be made. The authors provide a descriptive model of slope ratio estimation as a first step towards fully understanding the functions to consider when selecting an aspect ratio. Knowledge of how aspect ratios impact readability of different charts is relevant to our work, as some of our placement strategies are based on defining none, just one, or both of width and height of a word-scale visualization.

Finally, color perception is known to be impacted by changes in size [27] and, thus, is an additional important aspect to keep in mind when designing word-scale visualizations.

3 A DESIGN SPACE FOR WORD-SCALE VISUALIZATIONS

According to Schulz et al. [25] a design space consists of a finite number of design dimensions. Each of these dimensions captures one particular design decision needed to fully specify the whole. The dimensions of our design space describe constraints placed on the word-scale visualization design and text layout. Our goal is to provide a design-centered view on word-scale visualization placement that takes into account both the size of the visualization and the degree to which the visualization can disrupt the text.

3.1 Definitions and Framing

We first define word-scale visualizations, introduce important terminology used in our descriptions, and situate the design space within a larger set of layout choices.

3.1.1 Defining Word-Scale Visualization

We follow Tufte's general definition of sparklines (see Section 2.1) but loosen the strict "word-sized" requirement and consider a wider range of visualizations that are "word-scale"—usually larger than the size of a letter, but smaller than a sentence or paragraph. For example, the height of a word-scale visualization can be restricted to font height (as in most of Tufte's examples) but can also be larger. Word-scale visualizations can also overlap or fall within the *inter-line spacing*—the vertical white space between the bottom of one line of text and the top of the next—and/or the *inter-word spacing*—the horizontal white space between two consecutive words (see Figure 1). Allowing word-scale visualizations that are taller and wider than words introduces a number of alternative layout options, including using the existing inter-line and inter-word space, allowing visualizations to overlap with text, or changing the text layout to introduce additional space.

3.1.2 Static vs. Dynamic Integration

To place word-scale visualizations into text documents we have two fundamental options: static or dynamic integration. This is similar to




Fig. 1: Important terminology used in the paper in the context of entity, word-scale visualization, and surrounding text.

the labeling problem as discussed in Section 2.3.

With *dynamic integration*, word-scale visualizations are displayed in response to user-input. In the most limited case, a user might hover the mouse cursor over an entity revealing a small visualization next to it. Alternatively, user interaction could be used to simultaneously reveal or hide all visualizations in the document. The dynamic presentation of word-scale visualizations has the advantage that the text is not cluttered and can be read without the visualizations present. However, it also has several disadvantages: it requires an interactive environment and is not amenable to print media, showing visualizations individually can make comparisons more difficult, and—depending on placement options chosen—the text may have to reflow whenever visualizations appear or disappear.

Static integration on the other hand, means that visualizations are always present within the text. This strategy is well-suited to print media, and may also make it easier for users to compare visualizations (as noted by Zellweger et al. [34]). However, large numbers of word-scale visualizations might overwhelm or disturb reading the text.

Both static and dynamic integration entail trade-offs and it is up to the designer to make the right choice for their usage context. Our design space applies equally to static and dynamic placements but we do not discuss interaction techniques to invoke a dynamic placement.

3.1.3 Placement Context

We define three main placement options that relate to the contextual integration of word-scale visualizations in text. According to our definition, context is based on the relationship between an entity's bounding box and the bounding box of its associated visualization. The three types of contextual placement are:

Strong context: The entity's and visualization's bounding boxes touch or overlap—e. g. a word-scale visualization is placed just above or next to a word.

Weaker context: The two bounding boxes do not touch, but the position of the visualization is still defined by the entity's position in the text—e. g. a word-scale visualization is placed in the same general text region as the entity, but the two do not touch.

Out-of-context: Visualization position is not related to the entity position—e. g. a word-scale visualization is placed somewhere in the margin of the text.

Our design space concentrates on the strong context case where word-scale visualizations are placed in a sense "as close as possible" to the entity. We do not consider the weaker context and the out-of-context case as these introduce their own set of challenges related to linking the visualizations back to the entities. This is interesting area of study in which there has already been some promising research (e. g. Steinberger et al.'s *context-preserving visual links* [26]).

3.2 Design Space Dimensions

The dimensions for our design space can be roughly divided into two categories: dimensions that have an impact on the word-scale visualization (1–4) and dimensions that impact the text (5 and 6):

Control over the maximum height of the visualization: The height of a visualization is closely related to the amount of information encoded in its vertical dimension. A designer may have limited control over the height of a visualization if it is intrinsically bound to a property of the text, such as the font height or the inter-line spacing. Full control over visualization height means that the designer can freely choose any desired height and can make height choices based on the visual encoding.

Control over the maximum width of the visualization: The width of a visualization is closely related to the amount of information encoded in its horizontal dimension. A designer may have limited control over the width of a visualization if it is bound to properties of the text, such as the width of the entity. Full control over visualization width means that the designer can choose the width freely.




Fig. 2: Visual description of the three placement positions: *baseline* position, *top* position, and *right* position.

Word-scale visualization position: The strong-context position of a word-scale visualization is defined by a reference point on the visualization's bounding box—we use the left bottom corner throughout the remainder of the paper. Generally, the visualization can be placed anywhere around the entity as long as the two bounding boxes intersect. For simplicity we focus on three main positions:

- a) *baseline position*: on the left bottom corner of the entity's bounding box,
- b) *top position*: on the upper left corner of the entity's bounding box,
- c) *right position*: at the bottom right corner of the entity's bounding box (see Figure 2).

Visual encoding: The choice of visual encoding (e. g. line chart, bar chart, etc.) for the word-scale visualization has an effect on how much information can potentially be encoded within a given aspect ratio. The choice of data and visual encoding can thus effect how much control over width and height is necessary in order to ensure effective presentation of the data.

Amount of inter-word spacing introduced: Given a specific visualization width (freely chosen or defined through text properties), a designer can introduce additional inter-word space before or after the entity in order to control the amount of overlap between visualization and text.

Amount of inter-line spacing introduced: Given a specific visualization height (freely chosen or defined through text properties), a designer can decide to introduce additional inter-line space above or below the entity in order to control the amount of overlap between visualization and text. For example, additional inter-line space could be introduced to fit a visualization of given height between two lines of text without any overlap.

4 WORD-SCALE VISUALIZATION PLACEMENT IN PRACTICE

When considering how to place word-scale visualizations a designer may face several practical problems based on the characteristics of the document, the visualization, the data, and the usage scenario:

- **Is the text static, or can it be reflowed?** If word-scale visualizations are added to a scanned document, for example, it may not be possible to adjust the positions of words or phrases to accommodate the visualizations. However, in an electronic reading environment reflowing is easier and the inter-word spacing can be expanded to make space for word-scale visualizations.
- **Can the inter-line spacing be modified?** By increasing the space between lines, a designer can insert larger word-scale visualizations positioned above or below the text. However, adding inter-line spacing can also considerably increase the overall length and amount of whitespace in the document.
- **Should the visualizations be read along with the text?** If word-scale visualizations are intended to be read in the context of the sentences (as with the stock trends in Tufte's original example [29]), the designer may wish to place them in-line with the text. If the word-scale visualizations provide supplemental information that could disrupt reading, positioning it in the inter-line space may be more appropriate.
- **How important are the visualization's size and dimensions for readability?** The designer may need to enforce a minimum size for visualizations, e. g. for maps that contain small marks.

Line charts and other slope-based visualizations may also need to be rendered using a particular aspect ratio in order to facilitate accurate reading and comparison as discussed in Section 2.4

- **Is there an appropriate visual encoding for the data?** Visualizations that include text labels, axes, or visually complex marks may not be effective at small scale, or may need to be re-interpreted to work within the available space.

Answering these questions can be challenging, in part, because each individual choice may impact the size and layout of the document, as well as the readability of the text and word-scale visualizations. Moreover, the relationships between these decisions can be complex and their severity may depend heavily on the characteristics of the source document. For example, assume the following two designs:

- (A) Visualizations are placed in the *baseline* position and rendered in front of their entities at 20% opacity, but are displayed at full opacity when hovered with the mouse.
- (B) Visualizations are always shown at full opacity and placed to the *top* or *right* of their entities.

In Case A the visual encoding makes overlapping text and visualizations less problematic, but also makes comparison tasks more difficult. In Case B, comparison is easier. However, if there is not enough inter-line or inter-word space, overlap between visualizations and text *will* impede reading. A designer can avoid overlap by increasing the inter-line and/or inter-word spacing, but this can change the size and layout of the document considerably. To help illustrate the impact of and interactions between these design decisions, we discuss seven common placement options (illustrated in Table 1) in greater detail.

Case 1: Traditional In-Line Placement: The most common placement strategy for word-scale visualizations is to increase the inter-word spacing before or after an entity and to insert the visualization in this space. This gives the designer control over the width of the visualization and avoids collisions between visualizations and text, but restricts the maximum height of the visualization to the font height plus the height of the inter-line space. Text may also need to be reflowed if adding visualizations increases line lengths past the page or column width, which may be problematic for print documents. In-line placement may also be undesirable if visualizations are hidden and revealed in response to user interaction, since hiding or revealing a visualization may change the layout of all subsequent text.








Case 2: Overlay Word-Scale Visualizations on Entities:

Alternatively, we can choose to draw visualizations directly in front of or behind the entity. Since we introduce no additional whitespace, the layout of the text is not affected. However, the designer has restricted control over the width and height of the visualizations, which are limited by the width of the entity and font height plus inter-line spacing, respectively. This can make comparisons between visualizations difficult and can hurt readability if the entities are very short. In this placement strategy visualizations *always* collide with the entities they annotate, thus interaction may be necessary in order to disambiguate the two. For example, we can draw the visualization with a high transparency by default, but increase the opacity when a user hovers over it.

Case 3: Using Existing Inter-Line Space: For some documents it is possible to avoid layout changes and collisions by making visualizations the same width as the entity and positioning them within the existing inter-line space. In contrast to Case 2, this approach supports visual comparison between visualizations, but because visualization widths vary, those comparisons may still be difficult. Additionally, because the strategy constrains the visualization height to the current inter-line spacing, it is only viable for documents that are already widely-spaced, or that use very compact visual encodings.

Case 4: Using Inter-Line & Increased Inter-Word Space: To improve comparison and guarantee consistent visualization widths, we

Table 1: Several word-scale visualization placements options and the design space decisions that produced them.

case	control over max width of vis.	control over max height of vis.	introduce inter-word spacing	adjust inter-line spacing	overlap: vis.+ entity	overlap: vis.+ surr. text	possible collision between vis.	interaction potentially necessary	example placement
1	✓	—	✓	—	—	—	—	—	<i>uis autem vel pum iriure dolor</i> m consetetur  <i>sit amet,</i>
2	—	—	—	—	✓	—	—	✓	<i>uis autem vel pum iriure dolor .</i> m  <i>sit amet, dolor :</i>
3	—	—	—	—	—	—	—	—	<i>uis autem vel pum iriure dolor .</i> m consetetur  <i>sit amet, dolor :</i>
4	✓	—	✓	—	—	—	—	—	<i>uis autem vel pum iriure dolor</i> m consetetur  <i>sit amet,</i> <small>horizontal padding</small>
5	✓	✓	✓	✓	—	—	—	—	<i>uis autem vel pum iriure dolor</i> m consetetur  <i>sit amet,</i> <small>horizontal padding</small> <small>vertical padding</small>
6	✓	—	—	—	—	—	✓	✓	<i>uis autem vel pum iriure dolor</i> m consetetur  <i>color .</i> <small>overlap zone</small>
7	—	✓	—	—	✓	✓	✓	✓	<i>uis autem vel pum iriure dolor .</i> m  <i>sit amet, dolor .</i>

can combine the approaches from Case 1 and Case 3—placing fixed-width word-scale visualizations in the existing inter-line space, and increasing the inter-word spacing after the entity to avoid collisions. However, like Case 1, this approach means that text will be reflowed.

Case 5: Increasing Inter-Line & Inter-Word Space: We can extend the approach used in Case 4 and also insert additional inter-line space to accommodate visualizations with a chosen height. This provides the designer with full control over width and height of the word-scale visualization, while still avoiding collisions. However, introducing padding requires text to be reflowed between lines and can increase the length of the document (especially if visualization width and height are large). When selecting word-scale visualization sizes, a designer must compromise between visualization readability and the level of text disruption. While we know of no previous work that has studied this trade-off in detail, we provide a series of disruption and readability measures in Section 6 that can be used to evaluate candidate designs and inform designers’ decisions.

Case 6: Allowing Visualizations To Overlap: All of the previous cases have excluded the possibility of overlap between word-scale visualizations. However, relaxing this constraint introduces a number of new placement options. For example, by adding control over visualization width to the approach in Case 3, we can guarantee visualizations with consistent dimensions. This approach may be a good alternative if a text only has very few, widely spaced entities with visualizations. However, collisions between visualizations can appear when short entities occur in close proximity to one another (as seen in Table 1). An interaction (such as hovering the mouse over one of the visualizations) may be necessary to resolve these kinds of collisions.

Case 7: Allowing Visualizations to Overlap Surrounding Text: Finally, while all of the previous examples have prohibited visual-

izations from overlapping surrounding text, this restriction can also be relaxed. Any of the prior cases can be extended by allowing the height or the width to extend beyond the padded boundaries of the current word or line. This relaxation makes it possible for some visualizations to collide with the surrounding text, as well as with other visualizations, and interaction may be required to disambiguate them. Thus, this placement option may be most suitable when larger widths and heights are necessary to make word-scale visualizations readable, and no alternative visualization technique is viable.

While illustrative, these seven cases represent only a few points in a much larger space of placement options. Each option can be further modified by additional design decisions not yet discussed here.

For example, a number of possible strategies exist for selecting maximum visualization widths. A designer might choose the word-scale visualization width based on the number of dimensions in the data, the length of the longest entity, by using perceptual guidelines for chart banking [15], or some other approach. The length could be chosen to be the same for all visualizations or chosen on an individual basis. In summary, the cases above are meant to outline the space and serve as starting points for selecting specific variations appropriate for a given setting, audience, viewing scenario, or reading task.

5 WORD-SCALE VISUALIZATION PLACEMENT EXAMPLES

In this section we describe three “real-world” examples that we created to illustrate different types of word-scale visualization embeddings. These examples are meant to demonstrate the diversity of potential applications of word-scale visualizations and spark creativity in their application to text documents.

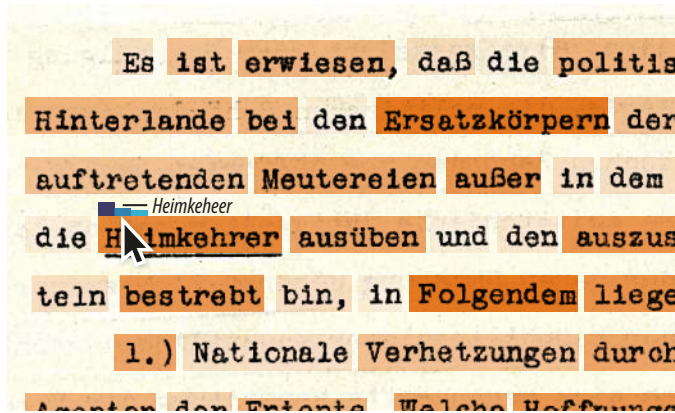


Fig. 3: Word-scale visualizations used in a scanned World War I artifact (courtesy of J. Benes). Words are overlaid with colored rectangles that express certainty of the OCR applied to the word. Small bar charts appear in the inter-line space when the mouse hovers over an entity showing alternative transcriptions of the word according to the OCR.

5.1 Word-Scale Visualization in Scanned Documents

Figure 3 shows a clipping of a scanned transcription of a historical World War I artifact from one of the historians we collaborate with. The text was scanned using a text recognition technique (OCR) to avoid manual transcription of the document. In this case, the historian was interested in knowing how exact the OCR technique was to potentially adjust the automatic transcription. In the specific example shown in Figure 3, every word became an entity as the OCR technique made a reasonable match for every word. To highlight uncertainty, we overlay rectangular bounding boxes on each entity and color-code them based on the uncertainty returned by the OCR algorithm (darker = more uncertain). When the mouse hovers over an entity, a tiny bar chart is drawn in the inter-line space, where each bar represents a possible alternative match. The height and color of each bar indicates the relative uncertainty of each of the alternative translations produced by the algorithm. Hovering over each bar shows the potential matching word and clicking a bar selects that match as the correct one.

In this example, one can consider the colored bounding box as one extremely simple word-scale visualization (a one-item heatmap) that has been placed at the baseline position and sized to the width and height of the entity. The small bar chart serves as a second word-scale visualization and is placed in the inter-line space. Together, the heatmap and bar chart illustrate two different placement strategies that are possible even if word positions cannot be adjusted. Additionally, this example highlights the importance of having a direct and close *strong context* association between visualization and entity. A viewer can look at each entity and decide whether the current match makes sense depending on available data, the semantic context, and syntax. Using the terminology from the previous section, this example illustrates both Case 3—because of the top position without overlap or control of visualization width and height—and a simple version of Case 2—due to the overlap of visualization and entity and the fixed size of the visualization. Interestingly, this particular version of Case 2 does not require interaction due to the simplicity of the visual encoding.

5.2 Word-Scale Visualization For Information Scent

Information scent is a viewer’s (imperfect) perception of the value, cost, or access path of information sources obtained from proximal cues [21]. Word-scale visualizations can be used to present proximal cues that provide information scent for viewers so they can make a decision whether or not it is worth their effort or time to pursue an exploration or reading path. Figure 4 shows a Wikipedia article augmented with word-scale visualizations placed above links to other pages. Each link in this example has become an entity paired with a horizontal-bar visualization whose length encodes the number of visits to the linked

Science fiction

From Wikipedia, the free encyclopedia

For other uses, see [Science fiction \(disambiguation\)](#).

Science fiction is a genre of [fiction](#) dealing with imaginative content such as [futuristic settings](#), [futuristic science](#) and [technology](#), [space travel](#), [time travel](#), [parallel universes](#), and [extraterrestrial life](#). It often explores the potential consequences of scientific and other innovations, and has been called a

Fig. 4: Small word-scale visualizations that provide information scent to help viewers decide whether or not they are interested in following an article linked to from a Wikipedia page. Small bar charts show how frequently the page linked to has been visited in the last 30 days. The word-scale visualizations were embedded with inter-line and inter-word spacing to prevent overlaps. Detail-on-demand from each data graphic gives information about the exact visit count for each link.

EASTERN EUROPE

Soviet cult and pragmatism in Transnistria

Experts worry that the next "Crimea" could be the breakaway region of [Transnistria](#). Many locals there don't share that view, and if the last referendum holds, a large majority would welcome a [Russian](#) annexation.

Fig. 5: Small maps embedded as word-scale visualizations provide details-on-demand for locations in a news story. Maps are embedded using a variant of the traditional (right) sparkline position. Interacting with a term or map expands it to show more detail.

page. We also draw a smaller gray reference line underneath each bar to aid comparisons between them. Here, we set the length of the reference line equal to the length of the bar for the most frequently visited page—in this case “futuristic” (which links to the Wikipedia page on “Future”). As in the previous example, interaction with the visualization is available. Hovering over the data graphic displays additional detail about the number of visits to the page in the last 30 days.

We set the height of the visualization equal to the height of the inter-line space and fixed the width of all visualizations at the average length of all entities in the text. In this example, we also added a small amount of inter-line space to reduce the level of visual clutter and ensure that the word-scale visualizations are associated with the entities below, rather than above them. To avoid overlap between visualizations we added inter-word spacing where the length of the visualization was longer than the length of its entity. We chose this position (Case 4) because the text contains a large number of entities and the traditional placement in the inter-word space (Case 1) might disrupt reading and would cause many consecutive words to be far apart.

5.3 Word-Scale Visualization for Detail-in-Context

Figure 5 shows a news story concerning the conflict between Russia and the Ukraine over Crimea. Within the news story the names of countries and regions are marked as entities. The word-scale visualizations for these entities are small geographical maps that are placed to the right of the word. The map visualizations provide location in-

formation that helps to put countries and regions in context. If a reader does not recognize a country from the small map, or wants more detail, they can hover over the visualization to reveal a larger version of the map with additional information. Here, we chose a variant of the Case 1 placement from Section 4 which introduces inter-word spacing and places each word-scale visualization slightly below the baseline position such that it is vertically centered with respect to the word. This slightly offset vertical position still avoids overlap between surrounding visualizations as long as the visualizations do not extend more than halfway into the inter-line space. In this example, adding word-scale visualizations required the text to be reflowed. However, because the number of regions mentioned in the text is limited and the map width is reasonably short, the amount of reflowing was minimal.

While it would be possible to add a larger map to the news article that shows all three regions in context, the word-scale visualizations have the benefit that they can a) show each region at reasonable size (both Crimea and Transnistria would become tiny when shown to scale next to Russia) and b) they save space in the article compared to a larger map in which all three regions would be visible. It would certainly be possible, however, to show the call-out of each map at a scale that included references to all three regions. For example, a call-out for Russia could show small pins to indicate where Transnistria and Crimea are located—even if they are too small to see.

6 EVALUATION

Decisions about how to integrate word-scale visualizations into a text document have two main effects: a) they constrain the layout of the text and the design of the visualizations, and b) they affect the legibility and understandability of the document as a whole. While both are important areas of study, we concentrate first on exploring how different word-scale visualization placement options alter the size and layout of documents. We conducted a quantitative evaluation in which we measured the impact of a number of different layout parameters, including visualization size and position, line height, and spacing. We ran the evaluation using real text documents and systematically varied the number of entities in the text so that we could characterize the trade-offs between placement options for both sparse and dense distributions of word-scale visualizations.

6.1 Documents

We randomly selected a set of 15 news articles from NLTK's ABC corpus of rural and science news [3]. These articles contained 100–246 words with average word lengths between 4.98–5.66 characters (median 5.29). We then created multiple versions of each article that contained varying numbers of entities. To create each new version, we automatically selected a number of words in the original text to mark as entities. We picked entities using a random log-normal distribution using a predefined average entity interval M (geometric mean) and 1.2 as the geometric standard deviation s , leading to $\mu = \log(M)$ and $\sigma = \log(s)$ as inputs to the log-normal random number generator. We generated tagged versions of each text using seven different entity intervals ($M=1, 5, 10, 15, 20, 25,$ and 30). This produced a final set of $15 \times 7 = 105$ articles with entity intervals ranging from 1 (one word on average between entities) to 30 (thirty words on average between entities, or about one entity every two lines). We then rendered these texts as HTML documents. To simulate the layout constraints of a common news or print article, we rendered the text using 12pt Times New Roman and set the maximum width of the document to 17cm, the equivalent of a standard A4 page with 2cm margins. We chose the font size based on work by Dyson et al. [9], who observed that 12pt type was read significantly more quickly than other type sizes.

6.2 Factors and Procedure

Using this set of HTML documents, we analyzed a series of different word-scale visualization placement options while varying five factors:

word-scale vis. position	<i>right / top / baseline</i>
word-scale vis. width	<i>entity len. / len. of shortest / len. of longest</i>
word-scale vis. height	<i>word height / 2.5×word height</i>
line spacing	<i>single-spaced / double-spaced</i>
spacing adjustment	<i>none / inter-line / inter-word / both</i>

The characteristic values for the line spacing, like font-size, were drawn from Dyson et al. [9]. For the visualization height we chose the word height (following Tufte's example) and also added the extreme case of $2.5 \times$ word height to observe a condition that exhibited vertical overlap even with double-spaced text (as in Case 7, Table 1). Combining these settings produced $3 \times 3 \times 2 \times 2 \times 4 = 144$ design alternatives, including examples illustrating each of the cases given in Section 4. We then used a sequence of scripts to automatically apply the parameters from each design alternative to our set of sample articles and measured the changes in word-scale visualization application and document layout.

6.3 Measures

To quantify how the different placements impacted the size and readability of the document, we computed four different metrics:

Change in Document Area: The increase in the total area of the document following visualization insertion, padding, and reflowing (if necessary), as a percentage of the original area.

Amount of Text/Visualization Overlap: The % of the surface area of the text in the document overlapped by word-scale visualizations. We count as overlapping area any place where the bounding box of one or more visualizations intersects the bounding box of a word.

Amount of Visualization Overlap: The % of the surface area of the visualizations overlapped by other word-scale visualizations. As before, we count as overlapping area any place where the bounding boxes of more than one visualization overlaps.

Text Shift: The average change in the X- and Y-positions of individual words in the document following word-scale visualization insertion, padding, and reflowing, as a percentage of the original document dimensions. Shifts in X occur when text is reflowed and moved within lines, while shifts in Y reflect both reflowing and changes to inter-line spacing. Since average shift in Y is linearly related to the change in document size, we do not report it independently.

6.4 Results

As part of our quantitative evaluation we computed placement statistics for each of the 144 different design alternatives using each of the 105 source documents for a total of $144 \times 105 = 15,120$ data points.

For the most part, the relationships between the different design dimensions are fairly predictable—as evidenced by the results in Figure 6, which highlights the results for each of our seven example cases. For example we see that Case 5—which introduces a considerable amount of vertical space to accommodate taller word-scale visualizations—causes the largest expansion of document area. Case 1—which adds horizontal padding and reflows the document by adding traditionally positioned visualizations, also adds space. Meanwhile, Case 4—which places word-scale visualizations in the existing inter-line space and adds horizontal space only to prevent overlaps on longer words—adds less total area. Because Case 1, Case 4, and Case 5 all add some horizontal space, all shift the average X-position of the words considerably, especially as the density of entities increases. However, the other four cases—which do not alter the inter-word or inter-line spacing—do not cause the text to shift at all. Instead, they increase the document size only where their visualizations overrun the original margins of the document.

Similarly, the placement options in Cases 2, 6, and 7—all of which allow overlaps—result in more frequent collisions between elements. In Cases 2 and 7, which place word-scale visualizations over the entity, the overlap between text and visualization increases fairly linearly with the number of entities (Figure 6, Visualization/Text Overlap). Meanwhile, Case 6—which allows overlap between adjacent

Word-Scale Visualization Parameters						Expansion of Document Area				Vis./Vis. Overlap			Vis./Text Overlap			Text Shift X				
Case	Word-Scale Vis. Position	Word-Scale Vis. Width	Word-Scale Vis. Height	Line Spacing	Spacing Adjustment	Entity Interval	0%	50%	100%	150%	0%	20%	40%	0%	20%	40%	60%	0%	10%	20%
							Case 1	right	longest entity length	word height	double-spaced	inter-word only	1	125.9%						
						5	25.1%											23.00%		
						30	4.8%											6.53%		
Case 2	baseline	entity length	word height	double-spaced	none	1								50.27%						
						5								13.40%						
						30								2.40%						
Case 3	top	entity length	word height	double-spaced	none	1	2.3%													
						5	2.3%													
						30	0.0%													
Case 4	top	longest entity length	word height	double-spaced	inter-word only	1	66.1%											23.07%		
						5	12.1%											15.60%		
						30	2.1%											2.13%		
Case 5	top	longest entity length	2.5 x word height	double-spaced	inter-line & inter-word	1	172.3%											23.07%		
						5	78.2%											15.60%		
						30	27.0%											2.13%		
Case 6	top	longest entity length	word height	double-spaced	none	1	11.1%				43.80%									
						5	4.4%				0.27%									
						30	0.3%				0.00%									
Case 7	baseline	entity length	2.5 x word height	double-spaced	none	1	3.2%				2.60%			64.47%						
						5	3.2%				0.53%			17.73%						
						30	0.0%				0.00%			3.33%						

Fig. 6: Parameters used to produce the seven placement example cases used in our evaluation, along with average change in document area, text/visualization and visualization/visualization overlap, and text shift statistics for each. For each case, the table includes values computed for text with varying entity intervals: 1 (one word on average between visualizations), 5 (on average five words between), and 30 (on average thirty words between). The bar chart overlays are normalized by column and colored by row. Cells with only zeros have been omitted.

charts—exhibits very few visualization collisions until the density of data graphics becomes high (Visualization/Visualization Overlap).

6.4.1 Comparing Vertical and Horizontal Space Adjustment

One interesting design trade-off appears when deciding where to insert space to accommodate a word-scale visualization. One can, for example, increase the inter-line space if the word-scale visualization does not fit above the entity, thus increasing the overall length of the document. However, if instead we add *inter-word* space and insert the visualization to the right of the entity, the text that follows it must be reflowed. Because word-scale visualization lengths are usually a small fraction of the line width, it is often possible to add multiple inter-word visualizations without increasing the length of the document. However, when a new line is necessary, the document must grow by the height of an entire line plus a full inter-line space.

The relative space-efficiency of these two placement strategies depends on several factors, including the amount of existing inter-line space and the frequency with which entities occur in the text. To explore the trade-off in greater detail, we ran a second experiment where we simulated placing word-scale visualizations into texts with varying inter-line spacings and numbers of entities. We fixed the visualization height to the font height and tested three common variations for visualization width (the width of the current entity, the width of the longest word, and the width of the average word). Based on observations from our larger study we focused on inter-line spacings between 1.5 (one-and-a-half-spaced) and 2.0 (double-spaced). We also used the same set of 15 source documents but focused on cases with higher numbers of entities (entity intervals $M = 1$ to $M = 10$, inclusive).

- word-scale vis. position** *right / top*
- word-scale vis. adjustment** *inter-line / inter-word*
- word-scale vis. width** *longest entity / average entity / entity*
- line spacing** *1.0 / 1.5 / 1.6 / 1.7 / 1.8 / 1.9 / 2.0*

Our results (Figure 7) highlight the relationship between the two

placement options, as well as inter-line space and word-scale visualization width. When the initial inter-line spacing is small, inserting visualizations to the right of entities can often considerably reduce the growth of the document—often by a factor of two or more. However, in cases where the inter-line spacing is already close to the visualization height and where the number of entities is high, positioning word-scale visualizations above the target becomes more efficient. This is especially true when longer visualizations are used, since these are more likely to trigger reflow events that add additional lines.

7 CONSIDERATIONS FOR WORD-SCALE VISUALIZATION DESIGN & PLACEMENT

Based on our theoretical and quantitative analyses of the design space, as well as our experience building our word-scale visualization placement test tool and examples, we discuss several considerations for designing and placing word-scale visualizations. We focus these considerations around three topics: placement advice, interaction, and word-scale visualization readability.

7.1 Placement Advice

One of the main problems when choosing a placement for word-scale visualizations is the wealth of inter-related factors. When the placement constraints are well-defined (as in our OCR example in Section 5.1), the space of options becomes relatively narrow and it is easier to make an informed decision. If there are no constraints, the options may become very overwhelming. In this case, we offer two general points of advice from our quantitative analysis:

Use inter-line space where available. If the inter-line space is sufficiently tall to accommodate the desired visual representations, placing them above the entity in the inter-line space is often the best choice. This placement strategy reduces the need to reflow or expand the text and—unless visualizations have very long aspect ratios—it typically results in little overlap. For some small visualizations like Horizon Graphs and small line charts Heer et al. [16] give experimentally-validated minimum sizes small enough to fit between 1.5- and 2.5-




Fig. 7: The effect of word-scale visualization placement strategy on the overall expansion of the document as more visualizations are added. Each small multiple illustrates the trade-off between inserting space for visualizations to the right of the entities they annotate (in orange) and inserting space for them above their entities (in green), for a particular combination of visualization width, inter-line spacing, and entity interval. Intersections between the lines highlight conditions where the relative space-efficiency of the two placement approaches changes.

spaced type. Unfortunately, similar readability or minimum size guidelines do not yet exist for most chart types, so the discrimination of whether the space is sufficient rests in the hands of the designer.

Placing visualizations to the right (usually) saves space. If text spacing must be modified, adjusting the inter-word spacing and placing visualizations adjacent to their entities is almost always more space-efficient than inflating the inter-line spacing to accommodate them (provided visualization heights and aspect ratios are word-like). As noted in Section 6.4.1, we found that adding additional inter-line space to position visualizations above their entity was more space-efficient only in a few cases—typically when the inter-line space was already very close to the desired height or when the number of entities in each line was very large. Right-side placement also has the added benefit that visualizations can typically be taller, including the height of the font plus whatever inter-line space is available.

7.2 Considering Interaction

Although we did not investigate options for interaction in depth, they should be considered for real-world applications of word-scale visualizations. We see four main reasons to include interaction: to resolve conflicts, to see visualizations on demand (hide and show), to increase saliency for individual visualizations when the initial state uses high transparency, and for details on demand:

Resolving conflict: When adjusting inter-line spacing and/or inter-word spacing is not possible due to constraints in the document, visualizations may overlap other visualizations or the surrounding text. In these cases interactions can be used to disambiguate them—for example, by bringing the one under the mouse cursor to the foreground.

Hide and show: In some cases, word-scale visualizations may not have to be constantly present. Using interactive controls, the reader can decide when to display them. This may be done at the document level, by including controls that reveal or hide all visualizations. Alternatively, visualizations can be revealed on a per-entity basis by clicking or hovering over the entity. An icon or other visual highlight can be used to suggest this possible interaction to readers.

Increasing Saliency: In order to reduce visual clutter, a designer may decide to draw all word-scale visualizations at medium or low opacity. To see a visualization more clearly, it can then be interactively activated (e. g. using hover) and drawn at full opacity.

Details on demand: Finally, word-scale visualizations may have interactions of their own similar to larger visualizations. In Section 5 we showed several detail-on-demand cases where hovering the mouse

over a visualization or even some of its individual components brought up additional information.

Of course, which interactions are most useful and how exactly the interaction should be implemented is highly context- and visualization-dependent. The exploration of word-scale visualization interactions is a rich avenue for future research.

8 SOFTWARE

As part of this research, we have built a general open-source library [13] that eases the process of integrating word-scale visualizations into HTML documents, and provides a range of options for adjusting the position, size, and spacing of visualizations within the text. The library includes default visualizations, including small line and bar charts, and can also be used to integrate custom word-scale visualizations created using web-based visualization toolkits such as D3 [5].

9 CONCLUSION

In this paper, we described a space of options for designing and placing word-scale visualizations in text. We focused, in particular, on how word-scale visualization placements constrain the design of visualizations and the layout of documents. We illustrated the space of alternatives by discussing seven common placement choices and described how they can be varied. We also showed three real-world examples that suggest how different placement options can be applied in practice and highlight the richness and diversity of possible word-scale visualization placements, applications, and designs. In order to measure the design-related impacts of word-scale visualization placements, we conducted a quantitative evaluation that measured the effect of different layout parameters, including visualization size, position, line height, and spacing on text layout. Finally, we proposed several general design considerations for placing word-scale visualizations. Based on the work presented here, we plan to design specific word-scale visualizations for augmenting a note-taking tool for historians. We also intend to evaluate the performance of different word-scale visualizations and placement options in this context. Finally, by providing tools and guidelines that make it easier to include word-scale visualizations in text, we hope to encourage other designers and researchers to explore new possible integrations of text and data.

ACKNOWLEDGMENTS

This work is sponsored by the French Research Organization, project grant ANR-11-JS02-003 and supported by the Collaborative European Digital Archive Infrastructure project CENDARI (cendari.eu).

REFERENCES

- [1] A. Abdul-Rahman, J. Lein, K. Coles, E. Maguire, M. Meyer, M. Wynne, C. R. Johnson, A. Trefethen, and M. Chen. Rule-based visual mappings—with a case study on poetry visualization. *Computer Graphics Forum*, 32(3):381–390, 2013.
- [2] E. Bertini, M. Rigamonti, and D. Lalanne. Extended excentric labeling. *Computer Graphics Forum*, 28(3):927–934, 2009.
- [3] S. Bird, E. Klein, and E. Loper. *Natural Language Processing with Python*. O'Reilly Media Inc., 2009.
- [4] R. Borgo, J. Kehler, D. H. Chung, E. Maguire, R. S. Laramée, H. Hauser, M. Ward, and M. Chen. Glyph-based visualization: Foundations, design guidelines, techniques and applications. In *Eurographics 2013—State of the Art Reports*, pages 39–63. The Eurographics Association, 2012.
- [5] M. Bostock, V. Ogievetsky, and J. Heer. D³ data-driven documents. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2301–2309, 2011.
- [6] U. Brandes, B. Nick, B. Rockstroh, and A. Steffen. Gestaltlines. *Computer Graphics Forum*, 32(3):171–180, 2013.
- [7] B.-W. Chang, J. D. Mackinlay, P. T. Zellweger, and T. Igarashi. A negotiation architecture for fluid documents. In *Proceedings of the Conference on User Interface Software and Technology (UIST)*, pages 123–132. ACM, 1998.
- [8] W. S. Cleveland, M. E. McGill, and R. McGill. The shape parameter of a two-variable graph. *Journal of the American Statistical Association*, 83(402):289–300, 1988.
- [9] M. C. Dyson. How physical text layout affects reading from screen. *Behaviour & Information Technology*, 23(6):377–393, 2004.
- [10] J.-D. Fekete and C. Plaisant. Excentric labeling: Dynamic neighborhood labeling for data visualization. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*, pages 512–519. ACM, 1999.
- [11] S. Few. Time on the horizon, Last read: March 2014. http://www.perceptualedge.com/articles/visual_business_intelligence/time_on_the_horizon.pdf.
- [12] J. Fuchs, F. Fischer, F. Mansmann, E. Bertini, and P. Isenberg. Evaluation of alternative glyph designs for time series data in a small multiple setting. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*, pages 3237–3246. ACM, 2013.
- [13] P. Goffin, W. Willett, J.-D. Fekete, and P. Isenberg. Sparklifierator, Last read: June 2014. <http://inria.github.io/sparklifierator/>.
- [14] B. Greenhill, M. Ward, and A. Sacks. The separation plot: A new visual method for evaluating the fit of binary models. *American Journal of Political Science*, 55(4):991–1002, 2011.
- [15] J. Heer and M. Agrawala. Multi-scale banking to 45 degrees. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):701–708, 2006.
- [16] J. Heer, N. Kong, and M. Agrawala. Sizing the horizon: The effects of chart size and layering on the graphical perception of time series visualizations. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*, pages 1303–1312. ACM, 2009.
- [17] M. R. Jakobsen and K. Hornbæk. Transient visualizations. In *Proceedings of the Conference on Computer-Human Interaction (OzCHI)*, pages 69–76. ACM, 2007.
- [18] B. Lee, N. H. Riche, A. K. Karlson, and S. Carpendale. SparkClouds: Visualizing trends in tag clouds. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):1182–1189, 2010.
- [19] J. Pearson, G. Buchanan, and H. Thimbleby. Improving annotations in digital documents. In *Research and Advanced Technology for Digital Libraries*, pages 429–432. Springer, 2009.
- [20] C. Perin, R. Vuillemot, and J.-D. Fekete. SoccerStories: A kick-off for visual soccer analysis. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2506–2515, 2013.
- [21] P. Pirolli and S. Card. Information foraging. *Psychological Review*, 106(4):643–675, 1999.
- [22] H. Reijner. The development of the horizon graph. In *Proceeding of Workshop From Theory to Practice: Design, Vision and Visualization Extended Abstracts of IEEE VisWeek*. Citeseer, 2008.
- [23] T. Ropinski, S. Oeltze, and B. Preim. Survey of glyph-based visualization techniques for spatial multivariate medical data. *Computers & Graphics*, 35(2):392–401, 2011.
- [24] T. Saito, H. N. Miyamura, M. Yamamoto, H. Saito, Y. Hoshiya, and T. Kaseda. Two-tone pseudo coloring: Compact visualization for one-dimensional data. In *Proceedings of the Conference on Information Visualization (InfoVis)*, pages 173–180. IEEE, 2005.
- [25] H.-J. Schulz, T. Nocke, M. Heitzler, and H. Schumann. A design space of visualization tasks. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2366–2375, 2013.
- [26] M. Steinberger, M. Waldner, M. Streit, A. Lex, and D. Schmalstieg. Context-preserving visual links. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2249–2258, 2011.
- [27] M. Stone. In color perception, size matters. *IEEE Computer Graphics and Applications*, 32(2):8–13, 2012.
- [28] J. Talbot, J. Gerth, and P. Hanrahan. An empirical model of slope ratio comparisons. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2613–2620, 2012.
- [29] E. R. Tufte. *Envisioning Information*. Graphics Press, Cheshire, CT, 1990.
- [30] E. R. Tufte. *Beautiful Evidence*. Graphics Press, Cheshire, CT, 2006.
- [31] M. O. Ward. A taxonomy of glyph placement strategies for multidimensional data visualization. *Information Visualization*, 1(3-4):194–210, 2002.
- [32] W. Willett, J. Heer, and M. Agrawala. Scented widgets: Improving navigation cues with embedded visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1129–1136, 2007.
- [33] D. Yoon, N. Chen, and F. Guimbretière. TextTearing: Opening white space for digital ink annotation. In *Proceedings of the Conference on User Interface Software and Technology (UIST)*, pages 107–112. ACM, 2013.
- [34] P. T. Zellweger, S. H. Regli, J. D. Mackinlay, and B.-W. Chang. The impact of fluid documents on reading and browsing: An observational study. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI)*, pages 249–256. ACM, 2000.